Mostofa Najmus Sakib

Contact No: (208) 570-9365 | Email: mostofa.najmus@gmail.com | Boise, ID, USA
github.com/Mostofa-Najmus-Sakib | linkedin.com/in/mostofa-sakib/

Core Technical Skills

Programming Languages: Python, Java, C#, SQL, JavaScript, C/C++, Shell, HTML, CSS, MATLAB, R
Generative AI & LLMs: Agentic AI, Retrieval-Augmented Generation (RAG), LangGraph, LangChain, Model Context Protocol (MCP), LangSmith, LlamaIndex, LLM Orchestration, Agent Workflow Design, Tool Calling, Prompt Engineering, Vector Databases, Embeddings, Semantic Search, Hybrid Search, Reranking, Document Ingestion & Retrieval, LLM Evaluation, Guardrails, LLM Observability
Machine Learning: Supervised/Unsupervised Learning, Classification, Regression, Clustering, Feature Engineering, Model Evaluation, Cross-Validation, Hyperparameter Tuning, Ensemble Learning, Recommender Systems, Anomaly Detection, Random Forest, XGBoost, PCA, SHAP, LIME, Scikit-learn
Data Engineering & Analytics: ETL/ELT Pipelines, Data Ingestion & Transformation, Data Processing, A/B Testing, Statistical Analysis, Exploratory Data Analysis, Data Visualization, Hypothesis Testing, Bayesian Statistics, PySpark, Kafka, Streamlit, Shiny, Shapely, Tableau, Microsoft Power BI, Snowflake, BigQuery, Redshift, Oracle, MSSQL, PostgreSQL, NoSQL
MLOps & Cloud: Azure, GCP, AWS, Databricks, Azure DevOps, MLflow, Docker, Kubernetes, Jenkins, Terraform, CI/CD, ONNX, TensorRT, Model Deployment, Model Monitoring & Observability, Model Retraining Automation, Git, GitHub, Bitbucket
Deep Learning & NLP: CNN, RNN, LSTM, GAN, Seq2Seq, Transformers, BERT, GPT, LLaMA, BERTopic, Encoder-Decoder Models, TensorFlow, Keras, PyTorch, Hugging Face, spaCy, NLTK, CoreNLP, Word2Vec, GloVe
Computer Vision: YOLOv5, YOLOv8, U-Net, Mask R-CNN, Faster R-CNN, 3D CNNs, OpenCV, MediaPipe
Software Development: Full-Stack Development, Backend Development, REST API Design, Microservices, Flask, FastAPI, Angular, Django REST Framework, API Integration, Database Design, Unit/Integration Testing, Agile Development, System Design

Work Experience

Data Science Intern | Simplot **May 2026 – Aug 2026**

- Engineered an enterprise **agentic RAG application** using LangChain, LangGraph, and Shiny, orchestrating retrieval and LLM-driven analysis across 15,000+ pesticide labels to recommend context-aware alternatives based on active ingredients, application constraints, location, and inventory availability.
- Designed and integrated Model Context Protocol (MCP)-based tools to connect AI agents with structured pesticide data and business systems, enabling standardized tool calling and context-aware retrieval within agentic workflows.
- Developed an LLM-powered label-comparison agent with LangSmith observability and guardrails to evaluate pesticide restrictions, application requirements, and product suitability, supporting traceable, context-grounded recommendations.
- Containerized backend and AI services with Docker and deployed them using Kubernetes.
- Partnered with domain experts and stakeholders to translate product-substitution requirements into data, retrieval, and analytics workflows, iteratively delivering functionality through Agile development.
- Managed feature development through sprint planning and Azure DevOps, converting business requirements into engineering tasks and prioritizing work across competing requirements.

Graduate Assistant | AI-based Security (AIbS) Lab, CS Department, Boise State University **Aug 2021 – Present**

- Developed reproducible LLM-based adversarial evaluation pipelines to assess the robustness of news believability models against linguistic manipulation, incorporating reinforcement learning-based defenses and automated experimentation workflows.
- Conducted computational social science research on the online information ecosystem, studying how linguistic, social, and AI-driven factors affect perceived credibility, news consumption, and information diffusion across digital platforms.
- Developed NLP and graph-learning models for news believability analysis using 14,535 user comments, 188 news articles, and believability judgments from 611 participants, improving classification accuracy from approximately **80%** to **87%** through believability-aware graph modeling.
- Architected and implemented a reusable NLP model that detects multiple sockpuppet accounts early on Wikipedia with just one edit, increasing detection efficiency by **5%**.

- Developed readability models for children using lexical, syntactic, and BERT-based features, improving accuracy and interpretability by approximately **10%** over baseline models.
- Built analytical dashboards and data products that transform large-scale research datasets into interpretable metrics, trends, and visualizations for technical and non-technical audiences.

AI/ML Intern (Mobility Analytics) | National Renewable Energy Laboratory (NREL) **May 2025 - Aug 2025**

- Built scalable data ingestion and analytics pipelines using Azure and Databricks to process 1M+ social-media interactions, applying LLM-based sentiment analysis to identify brand-level trends, temporal shifts, and actionable public-perception insights for autonomous-vehicle stakeholders.
- Collaborated with data engineers and research stakeholders to translate analytical requirements into reusable data and model-serving services.

Machine Learning and Data Analytics Intern | Idaho Transportation Department **May 2024 - Aug 2024**

- Applied YOLOv8 and U-Net-based computer vision models to detect asphalt pavement segregation from satellite, cellphone, and drone imagery, achieving 90% accuracy.
- Deployed a Flask-based ML framework on Google Cloud with automated ingestion, OCR text extraction, and notification workflows.
- Developed C# and Python backend components integrated with SQL Server to support data-processing and application workflows.

Graduate Research Assistant | Civil Engineering Department, BSU **Aug 2018 – July 2020**

- Developed an algorithm to differentiate between fraud and typing errors in construction datasets.
- Introduced a data-driven fraud detection pipeline and user interface for engineers and professionals to capture fraud incidents and overpayments on construction projects.

Freelance Full Stack Developer **Jan 2016 – July 2018**

- Developed full-stack web applications with frontend interfaces, REST APIs, backend services, and database integrations.

Education & Credentials

Ph.D. in Computing (Computer Science emphasis**) -** Boise State University (BSU) **Fall 2022 - Fall 2026**
Master's of Computer Science - BSU **Fall 2020 - Summer 2022**
Master of Science in Civil Engineering - BSU **Fall 2018 - Summer 2020**
Bachelor of Science in Civil and Environmental Engineering - IUT, Bangladesh **Jan 2012 - Dec 2015**

Relevant Projects

- **Credibility Drift Attacks on News Articles Using Adversarial LLM Transformations**: Developed an adversarial LLM-based rewriting system to study how linguistic cues, hedging, sensationalism, and rhetorical changes impact news believability and misinformation spread.
  Tech: Python, LLaMA, Adversarial NLP, Offline Reinforcement Learning, PyTorch, Docker, Git
- **Mobility Data Analytics Dashboard for Transportation Insights**: Built a transportation analytics system for AV and mobility datasets, including ETL pipelines, SQL-based processing, and routing-related insights.
  Tech: Python, FastAPI, PostgreSQL, Pandas, Plotly, GCP, Docker, Git
- **Financial Data Scraper + ML-Driven Stock Signal Platform**: Built an end-to-end financial analytics system with automated web scraping, time-series modeling, and ML-based buy-signal prediction exposed via APIs.
  Tech: Python, Django REST Framework, FastAPI, Selenium/Requests, PostgreSQL, NoSQL, Docker, CI/CD
- **Recommender System Using LLMs and Graph-Based User Modeling**: Developed a personalized recommendation engine using LLM embeddings, graph neural networks, and retrieval-augmented generation to deliver context-aware, behavior-driven recommendations.
  Tech: Python, Mistral embeddings, GNNs, RAG, FAISS, PyTorch, TensorFlow, Pandas, Flask/FastAPI, PostgreSQL, NoSQL, Docker, CI/CD, Git

Research Publications

1. Credibility Drift Attacks: LLM-Crafted Adversarial Manipulations That Flip News Believability - WSDM 2026
2. A Data Driven Approach to Public Opinion Mining on Autonomous Vehicles: Sentiment Analysis of Social Media Comments Using Large Language Models - TRB 2026
3. Understanding News Consumers' Perceptions of Believability: A Study of Real and Fake News - CSCW 2025
4. Opposites Attract? Ambivalence in Distinguishing Fake News and Predicting their Spread - ICWSM 2025
5. Automated Detection of Sockpuppet Accounts in Wikipedia - ASONAM 2022
6. Engage!: Co-designing Search Engine Result Pages to Foster Interactions - IDC 2021